


NASA/TM-2003-212421



# Team Software Development for Aerothermodynamic and Aerodynamic Analysis and Design

*N. Alexandrov, H. L. Atkins, K. L. Bibb, R. T. Biedron, M. H. Carpenter, P. A. Gnoffo,  
D. P. Hammond, W. T. Jones, W. L. Kleb, E. M. Lee-Rausch, E. J. Nielsen, M. A. Park,  
V. V. Raman, T. W. Roberts, J. L. Thomas, V. N. Vatsa, S. A. Viken, J. A. White, and  
W. A. Wood*  
*Langley Research Center, Hampton, Virginia*

---

November 2003

## The NASA STI Program Office ... in Profile

Since its founding, NASA has been dedicated to the advancement of aeronautics and space science. The NASA Scientific and Technical Information (STI) Program Office plays a key part in helping NASA maintain this important role.

The NASA STI Program Office is operated by Langley Research Center, the lead center for NASA's scientific and technical information. The NASA STI Program Office provides access to the NASA STI Database, the largest collection of aeronautical and space science STI in the world. The Program Office is also NASA's institutional mechanism for disseminating the results of its research and development activities. These results are published by NASA in the NASA STI Report Series, which includes the following report types:


- **TECHNICAL PUBLICATION.** Reports of completed research or a major significant phase of research that present the results of NASA programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA counterpart of peer-reviewed formal professional papers, but having less stringent limitations on manuscript length and extent of graphic presentations.
- **TECHNICAL MEMORANDUM.** Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.
- **CONTRACTOR REPORT.** Scientific and technical findings by NASA-sponsored contractors and grantees.
- **CONFERENCE PUBLICATION.** Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or co-sponsored by NASA.
- **SPECIAL PUBLICATION.** Scientific, technical, or historical information from NASA programs, projects, and missions, often concerned with subjects having substantial public interest.
- **TECHNICAL TRANSLATION.** English-language translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services that complement the STI Program Office's diverse offerings include creating custom thesauri, building customized databases, organizing and publishing research results ... even providing videos.

For more information about the NASA STI Program Office, see the following:

- Access the NASA STI Program Home Page at <http://www.sti.nasa.gov>
- E-mail your question via the Internet to [help@sti.nasa.gov](mailto:help@sti.nasa.gov)
- Fax your question to the NASA STI Help Desk at (301) 621-0134
- Phone the NASA STI Help Desk at (301) 621-0390
- Write to:  
NASA STI Help Desk  
NASA Center for AeroSpace Information  
7121 Standard Drive  
Hanover, MD 21076-1320

NASA/TM-2003-212421



# Team Software Development for Aerothermodynamic and Aerodynamic Analysis and Design

*N. Alexandrov, H. L. Atkins, K. L. Bibb, R. T. Biedron, M. H. Carpenter, P. A. Gnoffo,  
D. P. Hammond, W. T. Jones, W. L. Kleb, E. M. Lee-Rausch, E. J. Nielsen, M. A. Park,  
V. V. Raman, T. W. Roberts, J. L. Thomas, V. N. Vatsa, S. A. Viken, J. A. White, and  
W. A. Wood*  
*Langley Research Center, Hampton, Virginia*

National Aeronautics and  
Space Administration

Langley Research Center  
Hampton, Virginia 23681-2199

---

November 2003

## Acknowledgment

The authors would like to thank Charles Miller of the Aerothermodynamics Branch at NASA Langley Research Center, Hampton, Virginia, and Susan Hurd of NCI Information Systems, Inc., Hampton, Virginia, for enduring multiple manuscript reviews of this work.

The use of trademarks or names of manufacturers in this report is for accurate reporting and does not constitute an official endorsement, either expressed or implied, of such products or manufacturers by the National Aeronautics and Space Administration.

---

Available from:

NASA Center for AeroSpace Information (CASI)  
7121 Standard Drive  
Hanover, MD 21076-1320  
(301) 621-0390

National Technical Information Service (NTIS)  
5285 Port Royal Road  
Springfield, VA 22161-2171  
(703) 605-6000

# Abstract

A collaborative approach to software development is described. The approach employs the agile development techniques: project retrospectives, Scrum status meetings, and elements of Extreme Programming to efficiently develop a cohesive and extensible software suite. The software product under development is a fluid dynamics simulator for performing aerodynamic and aerothermodynamic analysis and design. The functionality of the software product is achieved both through the merging, with substantial rewrite, of separate legacy codes and the authorship of new routines. Examples of rapid implementation of new functionality demonstrate the benefits obtained with this agile software development process. The appendix contains a discussion of coding issues encountered while porting legacy FORTRAN 77 code to FORTRAN 95, software design principles, and a FORTRAN 95 coding standard.

# Introduction

The objective of the Fast Adaptive AeroSpace Tools (FAAST) program at NASA Langley Research Center is to develop fluid dynamic analysis and design tools (ref. 1). The four primary elements in FAAST are CAD-to-Grid Methods, High Energy Flow Solver Synthesis (HEFSS), Optimally Convergent Algorithms, and Efficient Adjoint Design Methods. This paper primarily focuses on the software development practices adopted by the HEFSS and design elements of FAAST.

HEFSS aims to develop an unstructured-grid flow solver for hypersonic flow applications. This solver is to have the same chemical, thermal, and turbulence modeling capabilities as are presently available in the Langley structured-grid flow solvers LAURA (ref. 2) and VULCAN (ref. 3). The switch to unstructured-grid technologies is seen as an enabling capability to efficiently handle complex geometries and is synergistic with the CAD-to-Grid Methods' unstructured-grid generation and adaptation efforts (refs. 4 and 5). This unstructured-grid hypersonic flow solver is to be obtained by extending the capabilities of the existing unstructured-grid transonic flow solver FUN3D (ref. 6). Factors related to the choice of FUN3D as the baseline unstructured-grid flow solver are discussed in appendix A on page 25.

The complexity of HEFSS exceeds that of prior fluid dynamic development projects at Langley and is the first tool developed by a sizable team.<sup>1</sup> The prior development of computational fluid dynamic (CFD) codes at Langley and at the former, co-located Institute for Computer Applications in Science

---

<sup>1</sup>See ONERA's elsA project (ref. 7) and DLR's MEGAFLOW project (ref. 8) for other examples of team CFD development.

Table 1. CFD Code, Architect, and Application Domain

CFL3D	Rumsey/Biedron	Structured-grid (SG) aerodynamics
LAURA	Gnoffo	SG hypersonic aerothermodynamics
VULCAN	White	SG hypersonic propulsion
TLNS3D	Vatsa	SG aerodynamics
OVERFLOW	Buning	Overset SG aerodynamics
USM3D	Frink	Unstructured-grid (UG) aerodynamics
NSU3D	Mavriplis	UG aerodynamics
FUN3D	Anderson	UG aerodynamics and design
FELISA	Peraire	UG hypersonic aerodynamics

and Engineering (ICASE) has consisted of one or two people working focused applications or algorithms. Even when more people contributed to the development of a code, one person contributed the bulk of the code and served as gatekeeper for any changes. Examples of this paradigm are listed in table 1. There is overlap in the capabilities of these codes because their development processes were all sufficiently rigid as to make it cheaper to develop an independent code with new functionality rather than to extend an existing code. HEFSS aims to break this cycle by developing an extensible product.

Heavyweight software engineering processes,<sup>2</sup> which accommodate teams of tens of hundreds of programmers working with a relatively well-defined set of requirements, were initially considered but rejected as being too restrictive for a research team of about 10 people. However, the emerging agile software development movement<sup>3</sup> is perceived to be well suited to the uncertain requirements and size of teams typically present in a research environment. The agile movement views software development as an empirical, rather than defined, process (ref. 9). To manage the empirical process, agile methods incorporate rapid feedback mechanisms to enable constant steering and place a renewed emphasis on the heart of software development—software craftsmanship (ref. 10).

Making the switch from a one-code, one-developer paradigm to a team-based approach is a significant culture change, but the ambitious goals of HEFSS provided a strong motivation to look past skepticism and overcome resistance to change. The scope of HEFSS required a group of developers (initially 12 people at 25 to 100 percent work levels) with diverse areas of expertise to collaborate on the software. The 18-month milestone for the project was to demonstrate the synthesis of the structured-grid physical models on a cylinder case by using an unstructured-grid discretization. In addition, the existing functionality of the baseline FUN3D code was to be maintained

---

<sup>2</sup>See [www.sei.cmu.edu/cmm](http://www.sei.cmu.edu/cmm) for example.

<sup>3</sup>See [www.agilealliance.org](http://www.agilealliance.org).

within the HEFSS code base. To compound matters, there was a lack of *team* software development expertise. This critical ignorance was overcome through consultant-led workshops, an invited lecturer series, a support contractor, and two team members aggressively pursuing software development best practices training (refs. 11 and 12).

This paper documents how the HEFSS team adapted and incorporated agile software development practices to develop the next generation CFD application software. No claims are made that the correct processes were chosen or that the current processes have fully matured. It is difficult to objectively gauge the HEFSS development process, but the project is ongoing, morale is high, its practices have been adopted by other teams, it was included in a group achievement award, and the local software engineering process group is using it as a model. The experience and lessons learned are offered as a case study, which may be useful to others with similar backgrounds and goals.

## Team Obstacles

Before embarking on a discussion of *team* software development, a synopsis of typical barriers to team *formation and viability* is necessary. Without providing a fertile team environment, the collaborative software development techniques laid out in the next section will not work.

According to reference 13, a true team can be identified by its low turnover, its strong sense of identity, its sense of eliteness, its joint ownership of the product, and its members deriving more enjoyment from their work than expected from the nature of the work itself. Unfortunately, while there is an array of team-building techniques available, there is no simple recipe for creating cohesive teams.

Meanwhile, methods known to destroy teams are well documented. Briefly, barriers to team formation and ongoing viability include these: lack of trust, the promotion system, defensive management, bureaucracy, physical separation, fragmentation of people's time, quality reduction of the product, phony deadlines, and clique control. For expanded discussion of these issues, consult references 13 and 14.

## Collaborative Software Development

CFD software development at Langley has traditionally been performed in a rather unconstrained, self-governed environment. As mentioned earlier, most codes have typically been developed by one, or perhaps two researchers. This paradigm has worked relatively well and has produced software packages widely used by industry and academia.

Unfortunately, such software development strategies often result in codes that are complex and burdensome to maintain, and frequently subsequent working groups produce distinct versions of the code that are often incompatible with each other and previously released versions. Moreover, cohesiveness and portability are typically lost, as additional researchers contribute to the code, using their own coding style and practices.

In contrast to this ad hoc approach to code development, the HEFSS team sought to incorporate the software industry’s best practices, not only because of the challenges of working as a cohesive team, but also to find methods which would extend the life cycle of the new code. Everyone on the team had experienced the pain of adding new capability to a large, existing code which was developed in an ad hoc manner. Even a seemingly innocuous bug fix was unnerving because there was no repeatable method to discover whether the fix would break existing capability in some subtle manner.

A survey of industry best practices for software development was conducted, which included sponsoring a local ICASE lecture series entitled “Modern Programming Practices.”<sup>4</sup> Meanwhile, two pathfinder projects were conducted to gain hands-on experience. Detailed discussion and extensive reference lists are available in references 11 and 12.

As described earlier, the emerging body of agile software development methodologies was determined to have the best fit with the inconstant nature of a scientific research environment. Specifically, Extreme Programming (XP) (ref. 15) appeared to be the most mature, although at the time, documentation was limited to a few websites.<sup>5</sup> In addition, recent experience with ISO 9001 edicts tended to steer the team away from defined process management techniques implicit in methodologies like the Capability Maturity Model<sup>®</sup> (ref. 16) and its associated Team Software Process<sup>SM</sup> (ref. 17).

The collection of collaborative software development practices described herein evolved from weekly meetings in which the challenges and possible solutions were discussed. Issues discussed cover fresh start versus retrofit versus restructuring of existing code, language selection, coding standards, modularization and maintainability versus efficiency, acceptance testing, source code management etiquette,<sup>6</sup> and documentation. As the HEFSS team initially struggled with and then embraced new software development practices, other teams (CAD-to-Grid, Design Optimization) within the FAAST project adopted many of the same practices.

Specific software development techniques are discussed in the following sections, namely: XP, project retrospectives, status meetings, other communication mechanisms, and documentation.

---

<sup>4</sup>See [www.icaselibrary.org/series/MPP](http://www.icaselibrary.org/series/MPP)

<sup>5</sup>[www.c2.com](http://www.c2.com) and [www.extremeprogramming.org](http://www.extremeprogramming.org).

<sup>6</sup>Source code management etiquette—when source code should and may be committed to a common repository.

## Extreme Programming

XP is founded on four values: communication, simplicity, feedback, and courage. It was designed to keep the right communications flowing by employing many practices that cannot be done without communicating. XP also gambles that it is better to do a simple thing today and pay a little more tomorrow for any necessary changes than to do a more complicated thing today that may never be used; that is, in this universe one cannot “save time.” Meanwhile, XP’s feedback mechanisms cover many time scales since optimism is an occupational hazard of programming and feedback is the treatment. Finally, courage enables one to escape local optima.

Built from this value system, XP consists of 12 practices shown in table 2. Also shown in the table is the level to which the HEFSS team has adopted each

Table 2. Current Level of XP Adoption

Practice	Adoption	Comments
Sustainable pace	Full	No compulsory overtime.
Metaphor	Full	Using naive metaphor, i.e., CFD jargon.
Coding standards	Full	See appendix B on page 31.
Collective ownership	Full	Anyone can change any piece of code.
Continuous integration	Full	Automated build and test on three computer architectures.
Small releases	Partial	A portion of code base is currently export restricted; seeking to relieve this constraint.
Test-driven development	Partial	FORTTRAN 90 unit test framework not widely used; however, Ruby codes are typically created using TDD.
Refactoring	Partial	Performed, but not mercilessly, due to lack of unit test coverage.
Simple design	Partial	Upfront, complex design is hard to resist, especially without strong test-driven development and refactoring.
Pair programming	Partial	Practiced, but not exclusively.
On-site customer	Partial	No outside customer is providing a business perspective, currently self-serving as customer for research products at hand.
Planning game	None	Have yet to invoke project management side of XP.

practice. The ensuing sections serve to briefly describe each practice and also to describe a practice in the context of the HEFSS team. Adjacent to the start of each section are quotations from reference 15.

## Sustainable Pace

Formally known as “40-hour week,” the sustainable pace practice probably ranks the highest on the common sense scale, but it is also the most frequently violated by managers and developers alike. Since the majority of the research conducted with the HEFSS project is years from commercial use, compulsory overtime is simply not part of the working environment.

Productivity does not increase with hours worked; tired programmers are less productive than well-rested ones.

Guide all development with a simple shared story of how the whole system works.

## Metaphor

Employing a system metaphor that all participants can understand facilitates communication both within the code and within the team. Since all the team members are familiar with CFD jargon, the naive metaphor is used.

Programmers write all code in accordance with rules emphasizing communication through the code.

## Coding Standards

Coding standards are usually dreaded and met with resistance because they are seen as adding a superfluous burden. After a brief discussion of the genesis of HEFSS’s coding standard, several reasons are provided to demonstrate why a coding standard is not only necessary but actually quite beneficial for a team software development project.

During the transition of legacy code from FORTRAN 77 to FORTRAN 95, a rough guess at a coding standard was created and used by the entire team. Based on this experience, a more detailed revision was created. (See appendix B on page 31.) One duty of the full-time contractor assigned to the team is to enforce the coding standard as new content is committed to the repository. This function is slated to be replaced by an automated agent that parses the source code.

Given a thoughtfully crafted coding standard, improved source code readability is a natural benefit through consistent indentation, alignment, naming, and commenting conventions. However, the coding standard must be appropriately tailored to the programming language. For example, FORTRAN 95 permits declaring an array variable and later dimensioning it through a separate statement. This multiline variable declaration can be hard to follow and can create confusion, thus prompting a line in the coding standard to place all attributes of the declaration on a single line, if possible. Another example is that the variable names of arguments in the calling and called routines do not have to match. However, retaining the same names for both improves global comprehension of the code and makes code-generated documentation more coherent.

A coding standard also serves as a sentinel against the use of vendor-specific language extensions or depreciated elements of the language that do not lend themselves to portability across various platforms. For example, FORTRAN 95 does not contain a complete set of intrinsic functions for accessing system-level

utilities or timing, but many compiler vendors offer extensions like `system()` and `etime()`, which are tempting but create portability headaches.

### Collective Ownership

Anyone can change any code anywhere in the system at any time.

The ideal situation for team software development occurs when a pair of developers looks at a given piece of code and does not feel the need to change the indentation, and so forth, and furthermore cannot recall whether they wrote the code in the first place. No single developer claims code ownership, yet all share responsibility; all source code is eligible for changes by any team member. Using a coding standard is absolutely essential to reach this goal.

Collective code ownership was a completely foreign concept to team members prior to this project. Initial acceptance of this philosophy came about because the original developer of the FUN2D/3D code was no longer at Langley, and the current “code steward” did not feel comfortable claiming the code as “his.” Both the software development practices mentioned above and the tools the team use for effective collaboration have cemented the idea of collective code ownership to the extent that members feel comfortable changing the code without asking permission of another developer.

Due to the team-oriented nature of the project and the amount of source code involved, a widely used source code management system is used, the Concurrent Versions System (CVS).<sup>7</sup> CVS oversees a central repository of the source code and allows each team member to concurrently develop and modify sections as needed. Any changes or additions to a local working copy can then be committed to the repository, whereby they will be available to the entire team.

CVS maintains complete documentation of any changes made during the course of code development, and previously modified or deleted code can be resurrected at any time by any member of the team. In addition, the system allows team members to work on platforms located virtually anywhere. The use of a software management tool allows for nearly seamless integration of a number of widely varying research projects and eliminates the need for multiple branches of a code.<sup>8</sup>

### Continuous Integration

Integrate and build the system many times a day, every time a task is completed.

In a team environment that has many developers who all contribute to a code base on a daily basis, integrating those changes into a common code base quickly becomes a major undertaking unless new code is integrated and tested as soon as practical, preferably within a few hours.

---

<sup>7</sup>[www.cvshome.org](http://www.cvshome.org)

<sup>8</sup>This CVS controlled L<sup>A</sup>T<sub>E</sub>X document was jointly composed by the team using such an approach.

Continuous integration avoids diverging or fragmented development efforts, in which developers are not communicating with each other about what can be shared or reused. Simply stated, *everyone* needs to work with the latest version of the code base. Making changes to obsolete code causes integration headaches.

Originally, developers manually ran the HEFSS test suite during code modification, but not all developers consistently ran the test suite before checking their code modifications into the repository, so an automated process was sought. At first the Unix-based `cron` utility was used to check out a fresh version of the CVS repository, to compile the suite of codes, and to run regression tests on three different architectures and compilers every night. However, the Extreme Programming community soon reminded the HEFSS team that “[daily builds] are for winning-challenged people who can’t integrate every 5 to 15 minutes and run all the tests at every integration,” and went to a true continuous integration mode of operation on dedicated machines.

The continuous integration process restarts the build and test process after each successful set of tests. Test results are automatically logged on a web server, and failures are e-mailed to all developers listing all CVS commits that were performed since the last successful build. With this system, errors are detected within a couple hours, and the integration failure e-mail provides a strong source of peer pressure on developers to run a range of tests before committing changes.<sup>9</sup>

### Small Releases

Put a simple system into production quickly, then release new versions on a very short cycle.

Feedback is the core idea behind the small releases practice. Get the software out there and learn from it. Strive to make the transition from pure software development to software maintenance as quickly as possible. Small releases are enabled by other practices like simple design, automated testing, and continuous integration.

The source code management system described previously enables the team to automatically create releases by merely “tagging” snapshots of the repository for which all the tests pass successfully during the continuous integration cycle. Routinely, the team typically makes several releases throughout any given day. This snapshot feature also facilitates the management of releases to outside users by providing accurate technical support tailored specifically to the exact source code snapshot released to a given party. Unfortunately, the HEFSS code currently has some restrictions on its external distribution; however, it is being used in-house by several people (ref. 18).

---

<sup>9</sup>See [www.martinfowler.com/articles/continuousIntegration.html](http://www.martinfowler.com/articles/continuousIntegration.html) for more information.

## Test-Driven Development

Since the time to fix a software defect (aka “bug”) scales exponentially with the time lag between introduction and detection (ref. 19), it is extremely advantageous to trap defects as early as possible during development.

Previously known merely as “Testing,” this practice has blossomed into a whole field in itself (ref. 20). Test-driven development within XP has two components, one centered around developers and the other centered around customers, or end users. Developers write unit tests so that their confidence in the code can become part of the code itself, while customers write acceptance tests so that their confidence in the code’s capabilities can also become part of the code. These automated tests allow confidence in the code to grow over time, allowing the code to become more capable of accepting change.

Unit tests are intended to verify small quanta of functionality within a code and should be automated and run to completion in fractions of a second. The unit tests serve as a development guide by specifying the desired capability, interfaces, and expected output of a functional unit. Unit tests also serve as mobility enablers during code architecture shifts to ensure a safe path was taken. Mobility allows code to be easily reused and to have functionality extended while safely maintaining current functionality. Note that in most cases, there will be more lines of unit test code than actual production code.

Acceptance tests check the interactions between code elements that unit tests cannot cover and document the existence of a particular code feature. Preferably, customers write acceptance tests.

Since the HEFSS code contains active research in many different disciplines that coexist in the same framework, work in one field can introduce errors in others through the common framework. These errors can go unnoticed if the code, in part and in whole, is not verified in a repeatable manner. One well-known approach to finding defects and ensuring that the code produces repeatable, verified answers is through automated testing.

For example, an unforeseen interaction with module A is introduced by modifying code in module B. If the problem in module A goes undetected for a month, it may be difficult to link the problem to an interaction with module B or to other code modifications made during that month. If the problem in module A is detected in minutes by an automated testing framework, the interaction of module A and module B can be clearly identified before other code modifications cloud the picture.

The current project began with legacy code that did not contain a single unit test. Because retrofitting an exhaustive set of unit tests to the existing legacy code was deemed too expensive, the original intent was to introduce unit tests as new code was added and old code was refactored. To date, however, unit testing has not been widely adopted by the team despite the creation of

Programmers continually write unit tests, which must run flawlessly for development to continue. Customers write tests demonstrating that features are finished. Any program feature without an automated test simply doesn't exist.

a unit testing framework for FORTRAN 95.<sup>10</sup> Currently, unit tests only cover a very small percentage of the code base. However, significant unit testing coverage is being built into Ruby-based wrappers used for testing and grid adaptation. Additionally, some of the low-level FORTRAN library routines are becoming test-infected, for example, character-to-number conversion routines and linear algebra routines.

The acceptance tests for the HEFSS code are a suite of over 240 regression tests performed by a series of Makefiles. These regression tests simply compare the convergence history of residual, force, and moment calculations (or other output appropriate to the code under test) to previously recorded executions to machine precision (not just 2–3 digits). These results are referred to as “golden files.” These test fixtures ensure that the current code gives the same discrete answer as the original golden file. Makefiles were initially selected to perform these tests because the tests were seen as a natural extension to code compilation.<sup>11</sup> The compile operations were incorporated into the tests, so the tests are always performed with an executable file produced from the current source files. Test cases can be run on an individual basis or as an entire suite.

The current set of acceptance tests for HEFSS was added incrementally to first cover the legacy functionality of FUN3D and then new functionality, as it was added to the suite. The Makefiles that perform the tests have become complex, hard to maintain, and are being replaced in an incremental fashion with unit-tested Ruby. This unit-tested Ruby framework should be much easier to maintain and allow more flexibility. The Ruby framework can be reused to link a number of the codes together to perform complex functions such as design optimization and grid adaptation, in addition to testing.

## Refactoring

Programmers  
restructure the  
system without  
changing its  
behavior to remove  
duplication,  
improve  
communication,  
simplify, or add  
flexibility.

To extend a code’s viable lifetime and strive for the simplest design that will work, developers need lots of practice modifying the design, so that when the time comes to change the system, they will not be afraid to try it. Constant refactoring is absolutely essential to keeping the cost-of-change curve from growing exponentially as time increases. Reference 21 teaches developers how to refactor and why.

Automated testing, as discussed earlier, is absolutely essential to refactoring. Without a safety net of tests, subtle shifts in the code’s fundamental architecture toward a more agile, clean, and understandable design is extremely difficult and frustrating. Testing allows developers to modify code that they did not write so that the original developer can be sure that modified routines

---

<sup>10</sup>To facilitate both the writing and running of unit tests for FORTRAN 95 source code, a testing framework called F95UNIT has been developed using Ruby. F95UNIT has a model similar to the unit-testing frameworks for other languages, e.g., JUnit, PyUnit, Ruby test/unit.

<sup>11</sup>If the code is modified and needs to be recompiled, it should also be tested.

still perform the original purpose *correctly*, if the appropriate unit tests pass. This process leads to an environment in which the tests are paramount and the code can be easily modified to add new functionality, improve speed, or become more readable.

Due in part to the lack of extensive unit testing in the HEFSS code, many refactorings are delayed, creating a backlog of work. Occasionally, the team will tackle some of these tasks, but so far the backlog continues to grow. A renewed effort at promoting the benefits of test-first programming is being made within the team by drawing attention to the inefficiencies inherent in the “Code-n-Fix” style of programming.

### Simple Design

Simple design is defined by two ideas: One is the YAGNI principle, otherwise known as, “you aren’t gonna need it,” and the other is a chant, “do the simplest thing that could possibly work.” These principles should be internalized and provide instinctive reactions to “gold plating” or other ideas that do not seem to fit the current task. A simple design should not contain ideas that are not used yet but that are expected to be used in the future. However, one should pay attention to the word “expected.” If you are somehow assured of the future and that a given idea will be necessary, design with it in mind, but do not implement it now because you will best know how to add it when the time comes.

As with refactoring, the lack of unit test coverage within HEFSS code makes this practice difficult to follow completely. For many developers, it is also typically contrary to years of prior practice; regardless, the team can now at least recognize complexity, and several major strides have been made to reduce existing manifestations.

### Pair Programming

The initial reaction to the idea of two people working on the same task at the same computer at the same time is usually negative. However, this reaction is typically caused by painful experiences associated with “pair debugging” or simply misunderstanding the true nature of pair programming itself. Pair programming is not one person programming while another person watches. It is more akin to an animated conversation, facilitated by a white board, where one participant might grab the marker from the other and make a change while the first is still talking. Pair programming should be highly dynamic, and the participants should be able to switch “driver” and “navigator” roles at any point. Besides making programming more fun, pair programming provides an extensive host of benefits, such as streamlining communication, propagating knowledge, and continuous code reviews. Pair programming also greatly enhances collective code ownership. For a detailed discussion of the art of pair

The system should be designed as simply as possible at any given moment; extra complexity is removed as soon as it is discovered.

All production code is written with two programmers at one machine.

programming, see reference 22.

Within the HEFSS team, frequent pair programming is highly encouraged but not mandated. Figure 1 shows an example of a dedicated pair program-



Figure 1. Pair programming station.

ming station which includes adjustable task chairs, an adjustable table, and multiple styles of wireless keyboards and mice. Note: the dual screens are attached to a single computer and simply provide more desktop space.<sup>12</sup> However, simply swapping a desk with a table is the essential step toward accommodating pair programming.

Pair programming is used for all aspects of code development, for example, debugging, teaching, refactoring, and adding new features. Intimately involving a number of researchers at the lowest levels of code development ensures a relatively high truck number.<sup>13</sup> Traditional CFD codes at Langley are developed by individuals or small teams and most of the resulting code base

---

<sup>12</sup>There is rumored to be a study which measured a 70 percent productivity increase for software developers by simply doubling the screen real estate.

<sup>13</sup>The truck number is the size of the smallest set of people in a project such that, if *all* of them got hit by a truck, the project would be in trouble. See [c2.com/cgi/wiki?TruckNumber](http://c2.com/cgi/wiki?TruckNumber) for further discussion.

has a truck number of 1 or perhaps 2, whereas the current collaborative team approach yields a value near 10.

### **On-Site Customer**

This XP practice is intended to remove the communication barriers present in a typical contracted piece of software where a slew of requirements and specifications are defined up front and then the “code monkeys” are let loose to grind out the required piece of software. The pitfalls with this sort of contract negotiation are many, the least of which is that the customers seldom know what they want before they see a working prototype. By placing an end user with the team, XP is nearly guaranteed of delivering a relevant, useful piece of software.

Include a real, live user on the team, available full-time to answer questions.

As discussed in reference 12, the scientific research environment often creates a situation in which the developers are their own customers. This scenario requires diligent role playing to keep technical and business needs separated. Currently, the HEFSS team members largely act as their own customers, with only very minor input from project stakeholders.

### **The Planning Game**

XP uses a four-dimensional space to plan and measure progress: time, cost, quality, and scope. Scope is typically ignored by many project-planning mechanisms, but it plays a central role in XP. The planning game has two levels: iteration planning and release planning. The basic premise of the planning game is that business people determine scope, priority, composition of releases, and dates of releases, while technical people provide estimates, design consequences, the process, and detailed scheduling.

Quickly determine the scope of the next release by combining business priorities and technical estimates; as reality overtakes the plan, update the plan.

As shown in table 2 on page 5, the HEFSS team has not yet begun using this practice. However, full-cost accounting practices now being put into place may force this final XP practice to be invoked.

### **Project Retrospectives**

Sometimes referred to as XP’s “thirteenth practice,” project retrospectives (ref. 23) are important components of tailoring a process to a given situation. Every few months, the team takes time to reflect on past events and accomplishments. The goal is not faultfinding but learning how to do better in the future. During these sessions, the team begins with a discussion guided by the following three questions: what has gone well, what could be improved, and with what new techniques or tools should the team investigate? Currently these sessions are not as formal or wide-reaching as some of the formats presented in reference 23.

## Scrum Status Meetings

A daily, stand-up meeting is normally associated with XP, but it is not explicitly called out as a practice or given much structure except that nobody can sit during the meeting; it should be short, and it should happen every day before developers start pair programming. The HEFSS team has adopted a similar, but more structured status meeting format from another agile methodology, Scrum (ref. 9).

A Scrum status meeting is held daily by an appointed “Scrum Master” and lasts no longer than 15 minutes. The meeting has an open attendance policy, but only team members are allowed to talk. The team members, in turn, succinctly report three things: what they *did* since the last meeting, what they *will do* by the next meeting, and what *got in the way* (impediments). Additional discussion during a Scrum is strictly limited to clarification-related questions and to note topics that will be discussed at a later time by interested parties. The Scrum master plays the role of gatekeeper and takes notes. Later, the Scrum master compares performance with past commitments and follows up on situations that appear to be stalled. Most importantly, the Scrum master is responsible for removing impediments.

Scrum status meetings have several benefits from a management perspective. They offer a quick and easy mechanism to collect data for status reports and yield an immediate sense of whether a team is in trouble. By using Scrums to their benefit, management can avoid what *Peopleware* (ref. 13) claims is the ultimate management sin: wasting people’s time.

Since the HEFSS team is currently dispersed throughout the local campus and most developers are not full time, the Scrum status meeting is only held weekly. In addition, the team also allots some time afterward to address any topics which may have arisen during the Scrum. This post-Scrum gathering is governed by Open Space’s Law of Two Feet,<sup>14</sup> which states that if during the course of any gathering, persons find themselves in a situation in which they are neither learning nor contributing, they must use their two feet and go to some more productive space.

## Other Communication Mechanisms

Since communication and cooperation are essential to the success of the effort, several additional tools are employed in addition to the communication mechanisms implicit in XP. The first is a Majordomo-based electronic mailing list which serves to facilitate communication among team members that are distributed across the local campus. In addition to the e-mail list and weekly meetings, the team also uses a web-based collaborative tool known as

---

<sup>14</sup>See [www.openspaceworld.com](http://www.openspaceworld.com) for more discussion.

a Wiki.<sup>15</sup> A Wiki allows users to freely create and edit web page content by using any web browser. Wikis have a simple text syntax for creating web page elements, and they dynamically create a new web page when they encounter a CamelCased word (a mixed-case word containing at least two capitals). The team uses the Wiki for a number of purposes. For example, the testing status page is contained in the Wiki so that anyone can add new data to the page. The Wiki is also used to share data for emerging test cases that have yet to be incorporated into the automated testing system, and it also serves as a repository for otherwise tacit knowledge, for example, CompilerNotes and CreatingNewTestCases.

## Documentation

Documentation for the HEFSS code takes many forms. While currently the HEFSS code itself lacks a formal user manual,<sup>16</sup> it does have a more exacting form of documentation, a large set of regression test cases. Each test case directory contains everything needed to run a given type of case and can usually be readily adapted to a new type of case.

Meanwhile, developers have three tools available for browsing the HEFSS code base. Code browsing can take many forms and be done for various reasons; consolidating them into a single tool has so far proven to be an elusive goal.

The simplest tool is a web-based rendering of the CVS repository, generated on-the-fly by the open source VIEWCVS<sup>17</sup> tool. This approach is based on the CVS repository's file directory structure and thus lacks the ability to navigate the source by using internal structure. However, it is the only tool that readily provides access to prior versions of the source code.

A second tool, developed by a support service contractor using C++, parses the source code and generates web-based output by using a commercial tool, UNDERSTAND FOR FORTRAN,<sup>18</sup> which extracts calling tree graphs and code statistics. The C++ code also creates tables of variable declarations and renders comments associated with routines that are placed according to the coding standard. The web pages generated by this tool include source code listings that have been formatted with line numbers and are keyword-colored to enhance readability.

A third code-browsing opportunity leverages the open source code documentation system, RDoc,<sup>19</sup> which was originally intended for documenting Ruby source. A short extension for this system was written to parse and

---

<sup>15</sup>[www.wiki.org](http://www.wiki.org)

<sup>16</sup>The user manual is being written.

<sup>17</sup>[viewcvs.sourceforge.net](http://viewcvs.sourceforge.net)

<sup>18</sup>[www.scitools.com/uf.html](http://www.scitools.com/uf.html)

<sup>19</sup>[rdoc.sourceforge.net](http://rdoc.sourceforge.net)

format FORTRAN 95<sup>20</sup> and has subsequently been accepted into the RDoc distribution. The RDoc system extracts a graph of the code source based on files, modules, and routines. From these data, it can generate frame-based web pages, XML, or Windows™ help files that can be used to navigate the calling structure.

## Experience Adding New Functionality

This section presents a few examples of new functionalities that have been incorporated into the code base. These additions have been facilitated by the current software development process. None of these extensions had been explicitly planned for during the initial code development, and the ease of their inclusion is testament to the agility of the process.

### Time-Accurate Simulations

In support of both passive and active flow control research at Langley, the perfect-gas capabilities in the solver have been extended to higher order temporal accuracy. The validity of the approach has been verified through numerical experiments in which an order property consistent with a second-order scheme has been demonstrated for turbulent flows. With a trivial amount of effort, the modifications required to obtain these results in the perfect-gas realm were extended to include reacting-gas simulations. Current work focuses on evaluating third- and fourth-order time-integration schemes for perfect-gas flows (ref. 24), which should also be readily extendable to more complicated physical models as needed.

### Incorporating Multiple Element Types

Initially, the HEFSS solver made sole use of tetrahedral element types to discretize a given domain. However, the ability to accommodate additional element types such as prisms, hexahedra, and pyramids provides greater flexibility to match a given element type to a particular flow topology, and the extension to include such elements in all aspects of the package is currently ongoing. This effort represents one, if not the, most substantial modifications to the software to date because it extends the fundamental data structure used throughout the code base. The pre-/post-processor and solvers, as well as all of their associated linearizations for optimization and adaptation, require considerable modification at the most fundamental levels. This undertaking has revealed many areas in which additional refactoring is still required before an acceptable level of modularity is achieved.

---

<sup>20</sup>This extension was accomplished with only 120 lines of code.

## Two-Dimensional Capability

A major advantage of pursuing mixed-element discretizations is the ability to recover a truly two-dimensional solution capability, which can be achieved through the use of prismatic and hexahedral elements in the spanwise direction, such that flux balances need only be performed in the plane of symmetry. Axisymmetric flows can also be readily accommodated by adding source terms. The benefits of such an approach are substantial, in that a separate code need not be maintained for such problems, a longtime burden for the original FUN2D/3D developers. In addition, all algorithms and physical models available in the three-dimensional path are immediately available for two-dimensional solutions, which allows basic research to be carried out on less costly two-dimensional problems. When computations are extended to three dimensions, the inconsistencies normally associated with switching between two separate solvers are no longer an issue, and the results are not contaminated by differences in discretizations or solution methods.

## Multigrid Algorithms

A major thrust of the FAAST project is aimed at achieving textbook multigrid efficiency (TME), an effort that could drastically reduce solution times for complex problems (ref. 25). Since the baseline unstructured-grid solver used as the foundation for the current work did not include options for multigrid acceleration, much work has focused on implementing such a capability.

The use of an agglomeration multigrid algorithm relies on an edge-based discretization of the governing equations; this requirement precludes the ability to compute solutions to the full Navier-Stokes equations on mixed-element grids. For this reason, a geometric non-nested multigrid approach has been initially chosen for the HEFSS solver. Operations such as coarse-grid partitioning and intergrid transfers in a complex domain-decomposed environment have been developed, and a multigrid algorithm has been implemented. Although this capability has been coded primarily with perfect-gas applications in mind, the scheme has been implemented such that users performing reacting-gas computations will also be able to make immediate use of this research without the need to duplicate extensive low-level code development typically associated with geometric multigrid on domain-decomposed unstructured meshes.

One component necessary to achieve TME is a line-implicit solver to overcome stiffness associated with high-aspect ratio grid elements. The ability to form lines suitable for implicit relaxation, to obtain an appropriate partitioning, and to perform an exact inversion along each line has been developed and is applicable to any set of physical equations being solved (ref. 26).

## Incorporating High-Energy Physics

The thermochemical nonequilibrium models in HEFSS are identical to those in LAURA, but their implementation is substantially different. LAURA made extensive use of precompiler directives that allocated memory and defined the code path according to a diverse set of options. This compilation strategy evolved from an absence of dynamic memory allocation capability in FORTRAN when LAURA was originally coded and because of a desire to completely eliminate any model-dependent conditional statements within loops that could compromise vector efficiency. Any change in the gas model required a recompilation of the source code. LAURA employs a script to guide a user through the various permutations and combinations of options, but the process is burdensome to a user conducting parametric studies. In contrast, HEFSS only needs to be compiled once on any platform, regardless of the desired physics model options.

Model parameters in LAURA are initialized in block data routines; these routines have been replaced by formatted data files that use conventional formatted reads and namelists in the HEFSS solver. Model parameters that are unlikely to be changed by the user (thermodynamic curve fit constants, species molecular weights, and heats of formation) are assembled in one set of data files. Gas model options that are likely to be changed by the user on a frequent basis, such as the chemical composition of the gases entering the domain or the thermochemical model, are assembled in a separate file. This separation minimizes the amount of setup required to perform a given analysis.

## Adjoint Solver and Sensitivity Analysis

As important as the software practices in this effort are to the development of new analysis capabilities, they are absolutely critical to the success of the design element under FFAST. In references [27](#), [28](#), [26](#), a discrete adjoint capability has been developed for the solver. This effort represents the only capability of its kind and relies on several hundred thousand lines of exact hand-differentiated linearizations of the preprocessor, flow solver, and mesh movement codes with respect to both the dependent variables and the grid coordinates. For free-stream conditions of Mach 0.84, a  $3.06^\circ$  angle of attack, and 5 million Reynolds number, sensitivity derivatives of the lift and drag coefficients, with respect to several shape design variables for fully turbulent flow over an ONERA M6 wing, (ref. [29](#)) that were computed by using the discrete adjoint formulation, are shown in table [3](#). The adjoint results are in excellent agreement with those obtained using a complex-variable approach (ref. [30](#)) with a step size of  $1 \times 10^{-30}$ . This accuracy can easily be compromised by a single error anywhere in the source code. With a dozen researchers modifying code on a daily basis, the use of continuous integration and automated testing is critical in maintaining such accuracy. Just as residual and force convergence

Table 3. Comparison of Discrete Adjoint and Complex Variable Design Variable Derivatives for Coefficients of Lift and Drag for Fully Turbulent Flow Over an ONERA M6 Wing

	Camber	Thickness	Twist	Shear	
$C_L$	0.956208938269467	-0.384940321071468	-0.010625997076936	-0.005505627646872	adjoint
	0.956208938269046	-0.384940321071742	-0.010625997076937	-0.005505627647001	complex
$C_D$	0.027595818243822	0.035539494383655	-0.000939653505699	-0.000389373578383	adjoint
	0.027595818243811	0.035539494383619	-0.000939653505699	-0.000389373578412	complex

histories are monitored to machine accuracy for the flow solver on several architectures, similar quantities are constantly tested for the adjoint solver and gradient evaluation codes. This constant testing ensures that discrete consistency between the analysis and design tools is always maintained, regardless of the modifications being implemented in other parts of the software.

Similar to the continuous integration and testing performed for the hand-differentiated code, a Ruby code has been developed similar to the effort described in reference 31 to automatically convert the codes in the HEFSS suite to a complex-variable formulation. This capability can immediately recover a forward mode of differentiation for the entire solver at any time, with no user intervention. This procedure is also continuously tested.

## Design Optimization

Approximation and Model Management Optimization (AMMO) techniques (refs. 32, 33, and 34) have been recently added to the HEFSS software set. AMMO is a methodology aimed at maximizing the use of low-fidelity models in iterative procedures with occasional but systematic recourse to higher fidelity models for monitoring the progress of the algorithm. In current demonstrations, AMMO has exhibited from three to five-fold savings in terms of high-fidelity simulations on aerodynamic optimization of 3D wings and multi-element airfoils, where simplified physics models (e.g., Euler) computed on coarse grids serve as low-fidelity models, while more accurate models (e.g., Navier-Stokes) computed on finer grids serve as high-fidelity models. AMMO was the first approach for using variable-fidelity models analytically guaranteed to converge to high-fidelity answers.

Because AMMO relies on using a variety of models in a single optimization run, maintaining continuous integration and consistency with the entire software set is especially crucial for obtaining stable optimization results. However, designing a testing strategy for optimization presents an interesting challenge because optimization algorithms require reasonably well converged analyses and are, therefore, expensive. Procedures for automated testing of optimization software is currently under development.

## Output Error Correction and Grid Adaptation

One thrust of the FFAST program is to develop a mathematically rigorous methodology to adapt a grid discretization to directly improve the calculation of an output function. An adjoint-based error correction and adaptation scheme has produced excellent results in two dimensions (ref. 35). This scheme is being extended to three dimensions and incorporated into HEFSS (ref. 5). This error correction and adaptation scheme requires the calculation of flow and adjoint residuals on embedded grids with interpolated solutions. The modularity of the HEFSS reconstruction, flux, and adjoint routines facilitated this calculation.

The interpolation of the solution onto the embedded grid requires the calculation of least-squares gradients. This gradient routine was readily shared between the flow and adjoint codes. The element-based interpolation scheme was developed test-first with the F95UNIT framework. The code to compute the flow and adjoint residuals consists of only a small driver routine; the remainder of the code is reused from the flow and adjoint solvers. The anisotropic adaptation metric is calculated with code that was also developed test-first by using the F95UNIT framework.

## Concluding Remarks

The Fast Adaptive AeroSpace Tools (FAAST) team uses techniques from the arena of commercial software development to implement an agile process for managing a development effort on a production software tool set. The agile aspects, such as collective ownership, simple design, and the lack of change boards, enable the rapid development of previously unplanned-for functionality. At the same time, the rigorous aspects, such as continuous integration and testing, maintain a stable base for the existing functionality.

An additional benefit of the present software process is that since there is only one code base for all of the development efforts, advances in one area of functionality immediately become part of the mainstream capability and are thus readily available to other researchers and users. For example, time-accuracy enhancements developed in the context of perfect gas flows were easily extended to apply for chemically reacting flows.

One remarkable aspect of this project is that developers who previously shuddered at the word “process” gelled into a team that uses a fairly rigorous, pervasive software process that they enjoy.

## Colophon

This paper is typeset in Donald Knuth’s Computer Modern Font with the *free*, multiplatform L<sup>A</sup>T<sub>E</sub>X<sup>21</sup> typesetting system using the NASA class developed by Wood and Kleb.<sup>22</sup> The auxiliary L<sup>A</sup>T<sub>E</sub>X packages are `array`, `dcolumn`, `fancyvrb`, `multirow`, `rcsinfo`, `tabularx`, `textcomp`, `time`, `url`, `varioref`, and `xspace`.

## References

1. Thomas, J. L.; Alexandrov, N.; Alter, S. J.; Atkins, H. L.; Bey, K. S.; Bibb, K. L.; Biedron, R. T.; Carpenter, M. H.; Cheatwood, F. M.; Drummond, P. J.; Gnoffo, P. A.; Jones, W. T.; Kleb, W. L.; Lee-Rausch, E. M.; Merski, N. R.; Mineck, R. E.; Nielsen, E. J.; Park, M. A.; Pirzadeh, S. Z.; Roberts, T. W.; Samareh, J. A.; Swanson, R. C.; Vatsa, V. N.; Weilmuenster, K. J.; White, J. A.; Wood, W. A.; and Yip, L. P.: *Opportunities for Breakthroughs in Large-Scale Computational Simulation and Design*. NASA/TM 2002-211747, 2002.
2. Cheatwood, F. M.; and Gnoffo, P.: *User’s Manual for the Langley Aerothermodynamic Upwind Relaxation Algorithm (LAURA)*. NASA TM-4674, 1996.
3. White, J. A.; and Morrison, J. H.: A Pseudo-Temporal Multi-Grid Relaxation Scheme for Solving the Parabolized Navier-Stokes Equations. AIAA Paper 99-3360, June 1999.
4. Jones, W. T.: GridExAn—An Integrated Grid Generation Package for CFD. AIAA Paper 2003-4129, June 2003.
5. Park, M. A.: Three-Dimensional Turbulent RANS Adjoint-Based Error Correction. AIAA Paper 2003-3849, June 2003.
6. Anderson, W. K.; Rausch, R. D.; and Bonhaus, D. L.: Implicit/Multigrid Algorithms for Incompressible Turbulent Flows on Unstructured Grids. *J. Comput. Phys.*, vol. 128, no. 2, 1996, pp. 391–408.
7. Cambier, L.; and Gazaix, M.: elsA: An Efficient Object-Oriented Solution to CFD Complexity. AIAA Paper 2002-0108, Jan. 2002.
8. Kroll, N.; Rossow, C. C.; Becher, K.; and Theile, F.: MEGAFLOW—A Numerical Flow Simulation System. ICAS 98-2.7.4, Sept. 1998.

---

<sup>21</sup>[www.ctan.org](http://www.ctan.org)

<sup>22</sup>Restricted availability through [lms.larc.nasa.gov/library](http://lms.larc.nasa.gov/library), or contact [Bill.Wood@NASA.Gov](mailto:Bill.Wood@NASA.Gov)

9. Schwaber, K.; and Beedle, M.: *Agile Software Development with Scrum*. Prentice Hall, Oct. 2001. See also <http://www.controlchaos.com>. Accessed 16 June 2003.
10. McBreen, P.: *Software Craftsmanship: The New Imperative*. Addison-Wesley, 2002.
11. Wood, W. A.; and Kleb, W. L.: Extreme Programming in a Research Environment. *Extreme Programming and Agile Methods—XP/Agile Universe 2002*, D. Wells and L. Williams, eds., Springer-Verlag, Chicago, IL, Aug. 2002, vol. 2418 of *Lecture Notes in Computer Science*, pp. 89–99.
12. Wood, W. A.; and Kleb, W. L.: Exploring XP for Scientific Research. *IEEE Software*, vol. 20, no. 3, 2003, pp. 30–36.
13. DeMarco, T.; and Lister, T.: *Peopleware: Productive Projects and Teams*. Dorset House, 2nd ed., 1999.
14. Lencioni, P.: *The Five Dysfunctions of a Team: A Leadership Fable*. Jossey-Bass, 2002.
15. Beck, K.: *Extreme Programming Explained*. Addison-Wesley, 2000.
16. Paulk, M. C.; Weber, C. V.; and Curtis, W.: *The Capability Maturity Model: Guidelines for Improving the Software Process*. Addison-Wesley, 1995.
17. Humphrey, W. S.; and Lovelace, M.: *Introduction to the Team Software Process*. Addison-Wesley, 1999.
18. Lee-Rausch, E. M.; Buning, P. G.; Morrison, J. H.; Park, M. A.; Rivers, S. M.; Rumsey, C. L.; and Mavriplis, D.: CFD Sensitivity Analysis of a Drag Prediction Workshop Wing/Body Transport Configuration. AIAA Paper 2003-3400, June 2003.
19. Boehm, B. W.: *Software Engineering Economics*. Prentice Hall, 1st ed., Oct. 1981.
20. Beck, K.: *Test Driven Development: By Example*. Addison-Wesley, 2002.
21. Fowler, M.: *Refactoring: Improving the Design of Existing Code*. Addison-Wesley, 1999.
22. Williams, L.; and Kessler, R.: *Pair Programming Illuminated*. Addison-Wesley, 2003.
23. Kerth, N.: *Project Retrospectives: A Handbook for Team Reviews*. Dorset House, 2002.

24. Carpenter, M. H.; Viken, S. A.; and Nielsen, E. J.: The Efficiency of High Order Temporal Schemes. AIAA Paper 2003-0086, Jan. 2003.
25. Thomas, J. L.; Diskin, B.; and Brandt, A.: Textbook Multigrid Efficiency for Fluid Simulations. *Annual Review of Fluid Mechanics*, vol. 35, 2003, pp. 317–340.
26. Nielsen, E. J.; Lu, J.; Park, M. A.; and Darmofal, D. L.: An Exact Dual Adjoint Solution Method for Turbulent Flows on Unstructured Grids. AIAA Paper 2003-0272, Jan. 2003.
27. Nielsen, E. J.: Aerodynamic Design Sensitivities on an Unstructured Mesh Using the Navier-Stokes Equations and a Discrete Adjoint Formulation. Ph.D. Thesis, Virginia Polytechnic Institute and State University, 1998.
28. Nielsen, E. J.; and Anderson, W. K.: Recent Improvements in Aerodynamic Design Optimization on Unstructured Meshes. *AIAA J.*, vol. 40, no. 6, 2002, pp. 1–9. See also AIAA Paper 01-0596.
29. Schmitt, V.; and Charpin, F.: Pressure Distributions on the ONERA-M6 Wing at Transonic Mach Numbers. AGARD AR-138, May 1979.
30. Anderson, W. K.; Newman, J. C.; Whitfield, D. L.; and Nielsen, E. J.: Sensitivity Analysis for the Navier-Stokes Equations on Unstructured Meshes Using Complex Variables. *AIAA J.*, vol. 39, no. 1, 2001, pp. 56–63. See also AIAA Paper 99-3294.
31. Martins, J. R. R. A.; Kroo, I. M.; and Alonso, J. J.: An Automated Method for Sensitivity Analysis Using Complex Variables. AIAA Paper 2000-0689, Jan. 2000.
32. Alexandrov, N. M.; and Lewis, R. M.: A Trust Region Framework for Managing Approximation Models in Engineering Optimization. AIAA Papers 96-4101 and 96-4102, Sept. 1996.
33. Alexandrov, N. M.; Nielsen, E. J.; Lewis, R. M.; and Anderson, W. K.: First-Order Model Management with Variable-Fidelity Physics Applied to Multi-Element Airfoil Optimization. AIAA Paper 2000-4886, Sept. 2000.
34. Alexandrov, N. M.; Lewis, R. M.; Gumbert, C. R.; Green, L. L.; and Newman, P. A.: Approximation and Model Management in Aerodynamic Optimization with Variable-Fidelity Models. *J. Aircr.*, vol. 38, no. 6, 2001, pp. 1093–1101.
35. Venditti, D. A.; and Darmofal, D. L.: Anisotropic Grid Adaptation for Functional Outputs: Application to Two-Dimensional Viscous Flows. *J. Comput. Phys.*, vol. 187, 2003, pp. 22–46.

36. Bibb, K. L.; Peraire, J.; and Riley, C. J.: Hypersonic Flow Computations on Unstructured Meshes. AIAA Paper 97-0625, Jan. 1997.
37. Venditti, D. A.; and Darmofal, D. L.: Adjoint Error Estimation and Grid Adaptation for Functional Outputs: Application to Quasi-One-Dimensional Flow. *J. Comput. Phys.*, vol. 164, 2000, pp. 204–227. See also AIAA Paper 99-3292.
38. Venditti, D. A.; and Darmofal, D. L.: Grid Adaptation for Functional Outputs: Application to Two-Dimensional Inviscid Flows. *J. Comput. Phys.*, vol. 176, 2002, pp. 40–69. See also AIAA Paper 2000-2244.
39. Venditti, D. A.: Grid Adaptation for Functional Outputs of Compressible Flow Simulations. Ph.D. Thesis, Massachusetts Institute of Technology, 2002.
40. Park, M. A.: Adjoint-Based, Three-Dimensional Error Prediction and Grid Adaptation. AIAA Paper 2002-3286, June 2002.
41. Matsumoto, Y.: *Ruby in a Nutshell*. O'Reilly, Sebastopol, CA, 2002.
42. Thomas, D.; and Hunt, A.: *Programming Ruby: The Pragmatic Programmer's Guide*. Addison-Wesley, 2001.

# Appendix A

## Code History and Architecture

This appendix provides a detailed review of the HEFSS code history and its architecture. The first section contains an explanation of how the baseline code, FUN3D, was selected, followed by a section describing how FORTRAN 95 was selected as the programming language. These reviews are followed by two sections that cover porting FUN3D to FORTRAN 95 and the use of object-oriented design principles.

### Baseline Code Selection

Three Langley unstructured-grid codes, USM3D, FUN3D, and FELISA (ref. 36), were considered as the initial template for the HEFSS code. FELISA, an inviscid, unstructured-grid flow solver, already has considerable success in the hypersonic domain. It also has equilibrium and thermochemical nonequilibrium gas models. While the addition of thermochemical nonequilibrium source terms, thermodynamic models, and transport models was perceived to be straightforward, considerable effort would have been required to introduce the viscous terms, the viscous flux Jacobians, and an implicit solution scheme. Both USM3D and FUN3D are highly successful codes for computing viscous flow on unstructured grids within the subsonic to low supersonic speed regimes. Ultimately, FUN3D was selected because it is more robust in the hypersonic domain, which is apparently attributable to its combination of Roe Flux Difference Splitting, flux reconstruction, and associated limiters. In addition, its discretizations are similar to LAURA, and the discrete adjoint capability for perfect gas design (refs. 26, 27, and 28) and grid adaptation (refs. 5, 35, 37, 38, 39, and 40) was judged particularly appealing for future hypersonic design and grid adaptation. A successful retrofitting of FUN2D with thermochemical nonequilibrium models confirmed the viability of this approach.

### Programming Language

Most of the CFD codes developed at Langley are written in FORTRAN 77 and often rely on nonportable extensions such as vendor-specific functions or links with C code. For the current project, the team sought a single, unifying standard language under which to develop new code. After surveying the available programming languages and deciding that a mixed-language code base would increase complexity too much, FORTRAN 95 was selected for the new suite of codes. FORTRAN 95 promises the numerical performance of FORTRAN 77 with the advanced features of other languages, such as dynamic memory allocation, derived types, recursion, and modules. This choice also allows a

relatively straightforward conversion of a substantial legacy code base written in FORTRAN 77.

The selection of FORTRAN 95 was tempered by the commitment to deliver a hypersonic flow simulation with thermochemical nonequilibrium on a geometrically simple configuration within 18 months. Adoption of a programming language significantly different from FORTRAN would have required a learning period for the majority of the team members, who were already proficient with FORTRAN 77. The time required to bring team members up to speed in a new language, plus the time required for conversion of legacy FORTRAN 77 to a language outside the FORTRAN family, was judged too costly, relative to the potential benefit offered by any other language.

FORTRAN 95 training was tailored to team needs in a two-part workshop. Dan Nagle, from Purple Sage Computing Solutions,<sup>A1</sup> spent a day with the team learning the HEFSS code objectives and the architecture of the legacy code. Using this material, he prepared a two-day course which highlighted FORTRAN 95 features suited to the HEFSS project.

Auxiliary scripting for controlling code compilation, templating, and testing is performed with Ruby (refs. 41 and 42) and Make.<sup>A2</sup> Ruby is an open source, object-oriented, threaded-scripting language with cross-platform support, while Make is an open source compilation tool.

## Porting and Restructuring

To lay a solid foundation for the new suite of solvers, FUN3D and the physical models from LAURA and VULCAN were ported from a mixture of C and FORTRAN 77 to FORTRAN 95. Porting FORTRAN 77 code to FORTRAN 95 was initially thought to be a simple process that could be accommodated by using a combination of homegrown scripts and a commercial software package, FORESYS<sup>™</sup>.<sup>A3</sup> FORESYS<sup>™</sup> was helpful when `implicit none` was requested because it would automatically declare all variables used in the routine. It also provided instructive diagnostics for various classes of errors during the conversion process and when replacing common blocks by modules. However, it invariably reformatted lines and destroyed symmetric forms of equations that had been carefully introduced by earlier authors, and it repositioned or silently eliminated comments. Eventually, Ruby and Perl scripts were crafted to handle tedious, error-prone operations such as code indentation and the conversion of continuation symbols without losing the comments and other structured formatting. The remainder of the conversion was done manually.

As the team had a chance to study the legacy structure, it became clear that

---

<sup>A1</sup>[users.erols.com/dnagle](http://users.erols.com/dnagle)

<sup>A2</sup>[www.gnu.org/software/make/make.html](http://www.gnu.org/software/make/make.html)

<sup>A3</sup>FORESYS<sup>™</sup> is a trademark of Connexite S.A., for more information see [www.simulog.fr/is/2fore1.htm](http://www.simulog.fr/is/2fore1.htm).

the old arrangements of common blocks and subroutines were counter to the modularity and extensibility the team was trying to create; so, during the port to FORTRAN 95, common routines and functions were extracted and placed in a single, shared library directory, while data structures such as boundary conditions, grid metrics, and solution quantities were generalized to handle an arbitrary number of equations and were encapsulated in derived types.

The use of derived types provides additional flexibility over FORTRAN 77; however, early versions of FORTRAN 95 compilers often displayed a significant performance penalty when these constructs were used in the computationally intensive regions of the solver.<sup>A4</sup> Consequently, the restructuring effort often required reworking these core routines to recover performance comparable to the legacy solver.

This transformation took nearly a year and was not without difficulties, but it was definitely a worthwhile effort because it gave team members hands-on experience with a code most had never seen before, instead of merely accepting the results of an automatic conversion. The conversion process also gave the team an opportunity to create and tailor a coding standard<sup>A5</sup> suited to their style and knowledge. In addition, the total lines of source code had been reduced by some 40 percent, in itself a significant benefit from the standpoint of code maintenance.

## Modularity and Encapsulation

Modularization, along with abstraction, information hiding, and encapsulation, are also means used to enhance code maintainability and bring the additional promises of code reuse, reduced complexity, extensibility, and orthogonality.<sup>A6</sup> Abstraction is the process of picking out common features of objects or procedures and replacing them with a single, more general function. Information hiding reduces complexity by hiding details of an object or function so the developer can focus on the object without worry about the hidden details. Encapsulation, or combining elements to create a larger entity, is one mechanism to achieve this goal.

The FORTRAN 95 constructs of modules, interface statements, public and private declarations, and derived types were employed to implement these ideas. FORTRAN 95 modules are similar to the class construct in object-oriented languages, while derived types are akin to structures. Modules were designed to abstract types of operations (e.g., file input/output, memory allocation, interprocessor communication, execution timing, linear algebra, and so on). Many modules employ a generic interface statement that automatically detects the type, kind, and rank of the calling arguments at compile time

---

<sup>A4</sup>See appendix C on page 34 for current results.

<sup>A5</sup>See appendix B on page 31.

<sup>A6</sup>In this case orthogonal is used in the sense of mutually independent or well separated.

and matches them to an appropriate low-level routine, which allows them to be largely independent of any particular flow solver since data are only exchanged through well-defined interfaces. Many of these FORTRAN 95 interface statements are produced automatically in the build process by a Ruby script which emulates the template system available in C++. In the remainder of this section, specific examples are given to demonstrate the benefits of modularization and data encapsulation.

### **Memory allocation**

Array memory allocation is handled by a single interface statement in a module that automatically detects the type, kind, and rank of the argument and calls the appropriate low-level routine for the allocation and initialization. This abstraction streamlines memory allocation requests throughout the code since memory tracking and diagnostics can be placed and maintained in a single location.

### **Parallel Communication**

Originally, the baseline solver relied on a shared-memory implementation specific to SGI® hardware and was not portable to the increasingly popular cluster-based, distributed-memory computing platforms. Moreover, the communication operations were dispersed throughout the solver, and any modifications to the communication model needed to be made in numerous locations throughout the code. In the current work, the message passing interface (MPI) standard was selected. Interprocessor communication has been abstracted from all but the lowest levels of the source code and is now encapsulated in a single module.

With this centralized approach to MPI communication, it is now trivial to make sweeping changes to the parallel aspects of the code, including completely removing it to produce a sequential version of the code. This abstraction also benefited the team when the high-energy, reacting-gas portion of the code was parallelized successfully on the first attempt. Normally, a developer would expect to spend considerable time debugging interprocessor communication.

### **Boundary Conditions**

Another area in which modularity and data encapsulation have provided a significant benefit is in the treatment of boundary conditions. The baseline FUN3D solver was extremely deficient in its ability to handle a wide range of boundary conditions. The user was restricted to inviscid, inflow/outflow, and viscous boundary types. Information required for these boundary types was contained in hard-coded data structures specific to each condition and were dispersed throughout the code. This design had become extremely limiting in

recent applications and was clearly not sufficient for extension to high-energy flows, where a large array of boundary condition types are required.

Using FORTRAN 95 derived types to encapsulate boundary condition information, the baseline solver was completely refactored to allow the straightforward addition of new boundary types. For any given boundary condition, all necessary data are contained in a boundary condition type. An array of these derived types then constitutes all boundaries in a given problem. For boundary conditions requiring additional physical data, a link to an additional data structure specific to that boundary condition is encapsulated. Derived types also allow the additional enrichment of the data structure without modifying argument lists. In this manner, any number of different boundary groups can be efficiently handled at the higher levels of the solver and unrolled for use as needed.

It should be noted that this data structure also allows for a natural handling of cost functions based on boundary data required for the design and grid adaptation capabilities within FAAST. Objective functions composed of viscous and/or pressure contributions can easily be specified on any subset or combination of boundary groups such that a specific flow feature or region of the domain can be targeted. For example, if it is determined that a strong shock on the outboard section of a wing is responsible for a severe wave drag penalty, a cost function can easily be formulated based solely on the contribution of that boundary group to the total drag. This method represents a substantial improvement over the baseline capabilities, in which all boundary groups necessarily contributed to a given cost function.

### **Gas Physics**

Modules, interfaces, and derived types are used extensively for the gas phase physics modules, which include thermodynamics, transport properties, thermal relaxation, and chemical kinetics. The thermodynamics module contains the initial interface from the flow solver to gas phase physics. The transport property module interfaces with the flow solver and the thermodynamics module to define molecular viscosity, conductivity, and species diffusivities. The thermal relaxation module is engaged when populations of excited states (rotational, vibrational, and electronic modes) cannot be defined by a single temperature. This module provides the source terms that define energy exchange among the available, thermally distinct modes. The chemical kinetics module provides source terms for the species continuity equations that define the rate of production or destruction of species.

In conclusion, it should also be noted that because the HEFSS project started with a large legacy code base and modularity and data encapsulation are elusive goals, which are really only earned through experience, code architecture changes are ongoing. In addition, there are drawbacks to modu-

larization that must be considered. For example, it was originally anticipated that compilers could optimize high-level constructs like derived types as if they were written using their lower-level counterparts. However, as appendix C on page 34 reveals, such is not always the case in practice.

# Appendix B

## Coding Standard

*Note: bracketed numbers refer to line numbers in the sample program which follows.*

### Style

- Free format with no character past column 80
- Indentation: begin in first column and recursively indent all subsequent blocks by two spaces.
- Start all comments within body of code in first column [42].
- Use all lowercase characters; however, mixed-case may be used in comments and strings.
- Align continuation ampersands within code blocks [77].
- No tab characters
- Name `ends` [85].

### Comments

- For cryptic variable names, state description using by a comment line immediately preceding declaration or on end of the declaration line [62].
- For subroutines, functions, and modules, insert a contiguous comment block immediately preceding declaration containing a brief overview followed by an optional detailed description [42].

### Variable Declarations

- Do not use FORTRAN intrinsic function names.
- Avoid multiline variable declarations.
- Declare `intent` on all dummy arguments [63].
- Declare the kind for all reals, including literal constants, using a kind definition module.
- Declare `dimension` attribute for all nonscalars [63].
- Line up attributes within variable declaration blocks.
- Any scalars used to define extent must be declared prior to use [60].
- Declare a variable name only once in a scope, including `use module` statements.

### Module Headers

- Declare `implicit none` [35].
- Include a public character parameter containing the CVS `$Id$` tag [37].
- Include a `private` statement and explicitly declare public attributes.

## Subroutines and Functions

- The first executable line should be `continue` [69].
- Use the `only` attribute on all `use` statements [58].
- Keep `use` statements local, i.e., not in the module header.
- Group all dummy argument declarations first, followed by local variable declarations.
- All subroutines and functions must be contained within a module.
- Any pointer passed to a subroutine or function must be allocated by at least size 1 to avoid null or undefined pointers.

## Control Constructs

- Name control constructs (e.g., `do`, `if`, `case`) which span a significant number of lines or form nested code blocks.
- No numbered do-loops.
- Name loops that contain `cycle` or `exit` statements.
- Use `cycle` or `exit` rather than `goto`.
- Use case statements with case defaults rather than if-constructs wherever possible.
- Use F90-style relational symbols, e.g., `>=` rather than `.ge.` [73].

## Miscellaneous

- In the interest of efficient execution, consider avoiding:
  - assumed-shape arrays
  - derived types in low-level computationally intensive numerics
  - `use` modules for large segments of data
- Remove unused variables.
- Do not use common blocks or includes.

## Illustrative Example

```
1  ! Define kinds to use for reals in one place
2
3  module kind_defs
4
5      implicit none
6
7      character (len=*), parameter :: kind_defs_cvs_id = &
8          '$Id: cs_example.f90,v 1.6 2003/12/03 14:11:51 kleb Exp $'
9
10     integer, parameter :: sp=selected_real_kind(P=6) ! single precision
11     integer, parameter :: dp=selected_real_kind(P=15) ! double precision
12
13 end module kind_defs
14
15 ! A token module for demonstration purposes
16
17 module some_other_module
18
19     implicit none
20
21     character (len=*), parameter :: some_other_module_cvs_id = &
22         '$Id: cs_example.f90,v 1.6 2003/12/03 14:11:51 kleb Exp $'
23
24     integer, parameter :: some_variable = 1
```

```

25
26 end module some_other_module
27
28 ! A collection of transformations which includes
29 ! stretches, rotations, and shearing. This comment
30 ! block will be associated with the module declaration
31 ! immediately following.
32
33 module transformations
34
35     implicit none
36
37     character (len=*), parameter :: transformations_module_cvs_id = &
38         'Id: cs_example.f90,v 1.6 2003/12/03 14:11:51 kleb Exp $'
39
40     contains
41
42     ! Computes a stretching transformation.
43     !
44     ! This stretching is accomplished by moving
45     ! things around and going into a lot of other details
46     ! which would be described here and possibly even
47     ! another "paragraph" following this.
48     !
49     ! This contiguous comment block will be associated with the
50     ! subroutine or function declaration immediately following.
51     ! It is intended to contain an initial section which gives
52     ! a one or two sentence overview followed by one or more
53     ! "paragraphs" which give a more detailed description.
54
55     subroutine stretch ( points, x, y, z )
56
57         use kind_defs
58         use some_other_module, only: some_variable
59
60         integer,          intent(in) :: points
61
62         ! component to be transformed
63         real(dp), dimension(points), intent(in) :: x, y
64         real(dp), dimension(points), intent(out) :: z ! transformation result
65
66         external positive
67         integer :: i
68
69         continue
70
71         i = 0
72
73         if ( x(1) > 0.0_dp ) then
74             call positive ( points, x, y, z )
75         else
76             do i = 1, points
77                 z(i) = x(i)*x(i) + 1.5_dp * ( real(i) + x(i) )**i &
78                     + ( y(i) * real(i) ) * ( x(i)**i + 2.0_dp ) &
79                     + 2.5_dp * real(i) + 148.2_dp * some_variable
80             enddo
81         endif
82
83     end subroutine stretch
84
85 end module transformations

```

# Appendix C

## Fortran 95 Considerations

The rationale for some elements of the coding standard presented in the previous section are discussed in this section.

### Best Practices

The use of `implicit none` minimizes the possibility of variable type errors. An example of a type error is when the implicit FORTRAN integer typing scheme creates integers for variable names beginning with the letters “i” through “n” when the user had intended a real variable. This unintended declaration type is avoided because `implicit none` requires every variable to be declared explicitly.

The use of `only`<sup>C1</sup> prevents unintended changes to values of other variables in the inherited modules. The `only` statement also facilitates finding the module that provides the inherited variable. To further restrict access to variables or subroutines in modules, a `private` statement is to be placed at the top of the module. An exclusive and explicit list of `public` entities is therefore required to share module data and methods outside the module. This exclusivity prevents unintended variable modifications.

Use of equality comparison with reals should be avoided because small, round-off errors may be present. The difference between the two variables is compared to an intrinsic function like `tiny()` to provide a more reliable comparison.

In general, the use of the `select case` conditional construct is more efficient than using an `if-elseif` construct since `if-elseif` might require several condition evaluations, while the `select case` only contains one condition evaluation. The `select case` construct is analogous to the deprecated `computed goto`.<sup>C2</sup> Also `select case` constructs convey control logic in clearer fashion and allow for cleaner error handling through the `default case`.

### Performance Considerations

Throughout the FORTRAN 95 restructuring of the FUN3D solver, several efficiency issues pertaining to advanced coding constructs were uncovered. Features such as derived types and modules are extremely attractive for communicating data; however, it was found that current FORTRAN 95 compilers often

---

<sup>C1</sup>For example, `use aModule, only : aVariable`

<sup>C2</sup>See [groups.google.com/groups?threadm=9o7uhi%24pus%241%40eising.k-net.dk](https://groups.google.com/groups?threadm=9o7uhi%24pus%241%40eising.k-net.dk) for further discussion.

failed to produce performance comparable to that of conventional FORTRAN 77 constructs such as passing data through calling argument lists.

### Data Sharing With Modules

An intermediate restructuring of FUN3D relied almost exclusively on the use of FORTRAN 95 modules. By eliminating virtually every argument list in the solver, an exceptionally clean code was obtained. However, in subsequent testing, this implementation was shown to be several times slower in execution speed than the legacy C/FORTRAN 77 solver. Upon closer inspection, it was found that the use of modules to communicate large segments of data can be extremely inefficient. To illustrate this degradation in performance, the test code included as appendix D on page 40 has been executed on a range of platforms and compilers as listed in table C1. Here, data are communicated with a

Table C1. Compilers Used in Performance Study

Vendor	Options	Release	O/S	Platform
Absoft™	-03 -cpu:p6	8.0-1	Linux® 2.4.18	Intel® P3
Compaq®	-arch ev67 -fast -04 -tune ev67	X1.1.1-1684	Linux® 2.4.2	Alpha EV67
HP®	-03	2.4	HP-UX® B.10.20	HP® 9000
IBM®	-05	7	AIX® 3	IBM® 7044
Intel®	-03 -ipo -wK	7.1-008	Linux® 2.4.18	Intel® P3
Lahey-Fujitsu	--o2 --nwarn -static --nsav --ntrace --nchk -x -	6.20a	Linux® 2.4.18	Intel® P3
NAG®	-04 -Wc,-malign-double -ieee=full -unsharedf95	4.2	Linux® 2.4.18	Intel® P3
NA Software	-fast	2.2-1	Linux® 2.4.18	Intel® P3
PGI®	-fast	4.1-1	Linux® 2.4.18	Intel® P3
SGI®	-02	7.3.1.2m	IRIX® 6.5	SGI® R10000
Sun <sup>SM</sup>	-fast	6.2-2	SunOS™ 5.8	Sun <sup>SM</sup> Blade1000

file I/O routine, as well as a routine that performs a large amount of arbitrary floating-point manipulations. In addition to an array A passed through a traditional argument list interface, an identical array B is also passed to and from the subroutines through the use of a FORTRAN 95 module. For this test, the extent of the arrays is 20M, a value on the order of that encountered in typical aerodynamic simulations. The results are normalized on the data obtained using the argument list model. As can be seen in table C2, use of the module

Table C2. Unformatted Disk I/O Using 20M Integers and 20M Reals

Compiler	Assumed size	Module	Derived type	Assumed shape
Absoft <sup>™</sup>	1.00	1.00	1.03	1.04
Compaq <sup>®</sup>	1.00	0.98	6.47	0.99
IBM <sup>®</sup>	1.00	1.03	1.01	1.03
Intel <sup>®</sup>	1.00	1.16	1.14	1.05
Lahey/Fujitsu	1.00	6.05	5.99	1.04
NAG <sup>®</sup>	1.00	1.02	1.22	1.03
NA Software	1.00	0.99	1.08	1.01
PGI <sup>®</sup>	1.00	0.98	1.00	0.98
SGI <sup>®</sup>	1.00	31.91	31.37	34.80
Sun <sup>SM</sup>	1.00	0.98	1.02	0.98

construct can incur severe penalties for unformatted disk I/O. The module interface is over 30 times slower than the data transferred via a conventional argument list on an SGI<sup>®</sup>. For floating-point arithmetic, the module interface exhibits run times on the order of 20 percent higher than the computations using data brought in through an argument list, as shown in table C3. Due

Table C3. Compute Work Using 20M Integers and 20M Reals

Compiler	Assumed size	Module	Derived type	Assumed shape
Absoft <sup>™</sup>	1.00	1.16	1.84	1.20
Compaq <sup>®</sup>	1.00	1.40	1.47	1.38
IBM <sup>®</sup>	1.00	2.76	2.76	2.76
Intel <sup>®</sup>	1.00	0.97	0.98	0.95
Lahey/Fujitsu	1.00	1.07	1.07	1.02
NAG <sup>®</sup>	Aborted			
NA Software	1.00	0.95	1.13	0.92
PGI <sup>®</sup>	1.00	1.96	1.96	0.94
SGI <sup>®</sup>	1.00	1.10	1.10	1.07
Sun <sup>SM</sup>	1.00	1.42	1.40	1.07

to this performance degradation, the module construct is employed sparingly in the HEFSS solver as a means to share large data structures. Only small amounts of data such as free-stream quantities, algorithmic parameters, and turbulence modeling constants are shared through modules.

### Derived Types

The baseline C/FORTRAN 77 solver was also refactored to make extensive use of the FORTRAN 95 derived type construct. The derived type is very attractive

in the sense that a number of related quantities can be encapsulated in a single variable, yielding relatively short argument lists throughout the code. Using this paradigm, variables related to the computational grid are stored in a `grid` type; solution-related variables are located in a `soln` type, and so forth. When a low-level routine requires a fundamental piece of data such as the coordinates of a grid point `i`, the information can be extracted as `grid%x(i)`, `grid%y(i)`, and `grid%z(i)`. Arrays of derived types are also supported under FORTRAN 95, making the implementation of algorithms such as multigrid and multiple instances of quantities, such as boundary groups, straightforward.

As in the case of modules, it was found that the use of derived types can also incur severe execution penalties. As shown in the last column of tables C2 and C3, a similar test to the one described previously has been performed on an array `C` transferred as the component of a derived type variable. It can be seen in tables C2 and C3 that this coding idiom can yield execution times more than 30 times slower for unformatted disk I/O and nearly a factor of three slower for floating-point operations over the argument list model.

The current HEFSS solver uses derived types to encapsulate much of its data structures; however, the components of these types required by low-level routines are extracted at the calling level and are received as conventional scalars and arrays in the I/O- and compute-intensive portions of the code. This model allows simple argument lists at the higher levels of the code, while maintaining the performance of the baseline solver. From a developer's point of view, derived types are one of the more useful enhancements of FORTRAN 95 over FORTRAN 77. They allow the developer to string together variables in meaningful groups and treat them as a single entity when desired. The HEFSS code uses a number of derived types. For example, the grid derived type contains all the information needed for the specification of the discretized mesh—`x,y,z` values for each point in space, cell volumes, cell-face normals and areas, connectivity information, and so on. Any of this information is available with the simple construct `grid%variable`, e.g., `grid%x`. Derived types may also be concatenated, extending their usefulness. For example, the grid derived type in the HEFSS code encompasses a boundary condition derived type that contains all the necessary data to impose boundary conditions—the physical condition (e.g., solid wall), the locations of points on the boundary, surface normals, and so forth. In addition, the definition of the derived type may be extended at a future date without affecting existing code. For example, adding a cell-face velocity for moving grid applications would involve a one-line addition to the type definition and would be completely transparent to sections of code not requiring this information.

### **Assumed-Shape Arrays**

As shown in tables C2 and C3 on page 36, some compilers treat arguments passed via assumed-shape arrays as poorly as they did derived types. Assumed-shape arrays can be noncontiguous, and thus interfacing to old FORTRAN 77 routines may require data to be copied to form a contiguous data block. These data copies can cause a large increase in the total memory required to compute a flow solution for some compilers as compared to others.

### **Memory Copies**

Occasionally, it is desirable to bring variables into a routine via argument lists rather than modules, as demonstrated in tables C2 and C3. However, unexpected behavior was detected on certain platform/compiler combinations when argument lists were combined with low-level module use. In these instances, the variables in the modules were not synchronized with the argument list variables. This synchronization issue was resolved when argument lists were used consistently throughout the subroutines that needed access to the data. It was eventually surmised that this problem was due to memory copies made by some compilers during a subroutine call. When that data copy was modified, it was no longer synchronized with the original data stored in the module and accessed with `use`. Also, on return from the subroutine, the local copy of the data was used to overwrite the data stored in the module, possibly erasing any modifications of the original data while the copy existed. This behavior appears to be very compiler and application specific and very difficult to detect and instrument.

### **Compilation Errors, Warnings, and Information**

The various compilers listed in table C3 generally have different sets of constructs that deem errors or produce a warning or other information. Some of the compilers are generally more lenient or particular than others when it comes to the constructs that are accepted as valid code for compilation. The FORTRAN 95 code base has benefited from exposure to a large number of different compilers. The coding standard contains guidelines for promoting portability. This portability experience was gained by exposure to multiple compilers, which makes it important to build and test on many different architectures/compiler, and which also results in a code base that is very portable.

### **Compiler Maturity**

In addition to the problems discussed with performance, errors have been found in a number of compilers. Some versions of the compilers have contained errors that have prevented them from successfully compiling HEFSS. Also, compiled

code will sometimes suffer run-time errors that are specific to the compiler or its version. Some compiler vendors have been very quick to respond to compiler bug reports, and others have ignored our requests for resolution of these errors.

# Appendix D

## Array Storage Performance Study Code

```

! $Id: test_array_storage_performance.F90,v 1.2 2003/09/08 20:34:47 atkins Exp $
!
! Preliminary stab at testing the relative performance of various
! array types: through argument lists as assumed-size, assumed shape,
! and derived type; and via module data.

module kind_definitions

  integer, parameter :: iKind = selected_int_kind(r=8)
  integer, parameter :: rKind = selected_real_kind(p=15)

end module kind_definitions

module module_data

  use kind_definitions, only: iKind, rKind

  implicit none
  integer(iKind), save :: module_length
  integer(iKind), dimension(:), pointer, save :: module_data_array_int
  real(rKind), dimension(:), pointer, save :: module_data_array_real

end module module_data

module type_definition

  use kind_definitions, only: iKind, rKind

  implicit none

  type derived
    integer(iKind) :: component_length
    integer(iKind), dimension(:), pointer :: component_array_int
    real(rKind), dimension(:), pointer :: component_array_real
  end type derived

end module type_definition

module test_various_array_types

  use kind_definitions, only: iKind, rKind
  use type_definition, only: derived
  use module_data, only: module_data_array_int, module_data_array_real, &
    module_length

  implicit none

  integer, save :: number_of_tests = 0
  integer, save :: assumed_size_time = 0
  integer, save :: assumed_shape_time = 0
  integer, save :: module_data_time = 0
  integer, save :: derived_type_time = 0
  integer, save :: deref_derived_type_time = 0

contains

  subroutine reset_counters()
    number_of_tests = 0
    assumed_size_time = 0
    assumed_shape_time = 0
    module_data_time = 0
    derived_type_time = 0
    deref_derived_type_time = 0
  end subroutine reset_counters

  subroutine read_array_types(logical_unit, array_size, &
    assumed_size_array_int, assumed_size_array_real, &
    assumed_shape_array_int, assumed_shape_array_real, &
    derived_type)

    integer, intent(in) :: logical_unit
    integer, intent(in) :: array_size

    integer(iKind), dimension(array_size), intent(inout):: assumed_size_array_int
    real(rKind), dimension(array_size), intent(inout):: assumed_size_array_real

```

```

integer(iKind), dimension(:), intent(inout) :: assumed_shape_array_int
real(rKind),    dimension(:), intent(inout) :: assumed_shape_array_real

type(derived), intent(inout) :: derived_type

integer(iKind), dimension(:), pointer      :: dummy_array_int
real(rKind),    dimension(:), pointer      :: dummy_array_real

integer :: i
integer :: start, finish
integer :: size

continue

number_of_tests = number_of_tests + 1

rewind(logical_unit)
call system_clock(start)
read(logical_unit) (assumed_size_array_int(i), i=1,array_size)
read(logical_unit) (assumed_size_array_real(i),i=1,array_size)
call system_clock(finish)
assumed_size_time = assumed_size_time + finish-start

rewind(logical_unit)
call system_clock(start)
read(logical_unit) (assumed_shape_array_int(i), i=1,array_size)
read(logical_unit) (assumed_shape_array_real(i),i=1,array_size)
call system_clock(finish)
assumed_shape_time = assumed_shape_time + finish-start

rewind(logical_unit)
call system_clock(start)
size = module_length
read(logical_unit) (module_data_array_int(i), i=1,size)
read(logical_unit) (module_data_array_real(i),i=1,size)
call system_clock(finish)
module_data_time = module_data_time + finish-start

rewind(logical_unit)
call system_clock(start)
size = derived_type%component_length
read(logical_unit) (derived_type%component_array_int(i), i=1,size)
read(logical_unit) (derived_type%component_array_real(i),i=1,size)
call system_clock(finish)
derived_type_time = derived_type_time + finish-start

rewind(logical_unit)
call system_clock(start)
size = derived_type%component_length
dummy_array_int => derived_type%component_array_int
dummy_array_real => derived_type%component_array_real
read(logical_unit) (dummy_array_int(i), i=1,size)
read(logical_unit) (dummy_array_real(i),i=1,size)
call system_clock(finish)
deref_derived_type_time = deref_derived_type_time + finish-start

!   rewind(logical_unit)
!   call system_clock(start)
!   read(logical_unit) (assumed_size_array_int(i), i=1,array_size)
!   read(logical_unit) (assumed_size_array_real(i),i=1,array_size)
!   call system_clock(finish)
!   assumed_size_time = assumed_size_time + finish-start

end subroutine read_array_types

subroutine work_with_array_types(array_size,          &
                                assumed_size_array_int, assumed_size_array_real, &
                                assumed_shape_array_int, assumed_shape_array_real, &
                                derived_type)

integer, intent(in) :: array_size

integer(iKind),dimension(array_size), intent(inout):: assumed_size_array_int
real(rKind),    dimension(array_size), intent(inout):: assumed_size_array_real

integer(iKind), dimension(:), intent(inout) :: assumed_shape_array_int
real(rKind),    dimension(:), intent(inout) :: assumed_shape_array_real

type(derived), intent(inout) :: derived_type

integer(iKind), dimension(:), pointer      :: dummy_array_int
real(rKind),    dimension(:), pointer      :: dummy_array_real

integer :: i, j
integer :: start, finish
integer :: size

real(rKind)                                :: dot, sum, max_element

```

```

real(rKind), dimension(:), allocatable :: vector

continue

number_of_tests = number_of_tests + 1

allocate( vector(array_size) )

call random_number( vector )

call system_clock(start)
do j = 1, 5
  dot = 0.0_rKind
  do i = 1, array_size
    dot = dot + assumed_size_array_real(i)*vector(i)
  end do
  max_element = -huge(max_element)
  do i = 1, array_size
    if ( abs(assumed_size_array_real(i)) > abs(max_element) ) &
      max_element = assumed_size_array_real(i)
  end do
  sum = 0.0_rKind
  do i = 1, array_size
    sum = sum + assumed_size_array_real(i)
  end do
  do i = 2, array_size
    assumed_size_array_real(i) = assumed_size_array_real(i) &
      + assumed_size_array_real(i-1)
  end do
  do i = 1, array_size
    assumed_size_array_real(i) = assumed_size_array_real(i) &
      + assumed_size_array_real(assumed_size_array_int(i))
  end do
end do
call system_clock(finish)
assumed_size_time = assumed_size_time + finish-start

call system_clock(start)
do j = 1, 5
  dot = 0.0_rKind
  do i = 1, array_size
    dot = dot + assumed_shape_array_real(i)*vector(i)
  end do
  max_element = -huge(max_element)
  do i = 1, array_size
    if ( abs(assumed_shape_array_real(i)) > abs(max_element) ) &
      max_element = assumed_shape_array_real(i)
  end do
  sum = 0.0_rKind
  do i = 1, array_size
    sum = sum + assumed_shape_array_real(i)
  end do
  do i = 2, array_size
    assumed_shape_array_real(i) = assumed_shape_array_real(i) &
      + assumed_shape_array_real(i-1)
  end do
  do i = 1, array_size
    assumed_shape_array_real(i) = assumed_shape_array_real(i) &
      + assumed_shape_array_real(assumed_shape_array_int(i))
  end do
end do
call system_clock(finish)
assumed_shape_time = assumed_shape_time + finish-start

call system_clock(start)
do j = 1, 5
  dot = 0.0_rKind
  do i = 1, array_size
    dot = dot + module_data_array_real(i)*vector(i)
  end do
  max_element = -huge(max_element)
  do i = 1, array_size
    if ( abs(module_data_array_real(i)) > abs(max_element) ) &
      max_element = module_data_array_real(i)
  end do
  sum = 0.0_rKind
  do i = 1, array_size
    sum = sum + module_data_array_real(i)
  end do
  do i = 2, array_size
    module_data_array_real(i) = module_data_array_real(i) &
      + module_data_array_real(i-1)
  end do
  do i = 1, array_size
    module_data_array_real(i) = module_data_array_real(i) &
      + module_data_array_real(module_data_array_int(i))
  end do
end do
end do

```

```

call system_clock(finish)
module_data_time = module_data_time + finish-start

call system_clock(start)
do j = 1, 5
  dot = 0.0_rKind
  do i = 1, array_size
    dot = dot + derived_type%component_array_real(i)*vector(i)
  end do
  max_element = -huge(max_element)
  do i = 1, array_size
    if ( abs(derived_type%component_array_real(i)) > abs(max_element) ) &
      max_element = derived_type%component_array_real(i)
  end do
  sum = 0.0_rKind
  do i = 1, array_size
    sum = sum + derived_type%component_array_real(i)
  end do
  do i = 2, array_size
    derived_type%component_array_real(i) &
      = derived_type%component_array_real(i) &
      + derived_type%component_array_real(i-1)
  end do
  do i = 1, array_size
    derived_type%component_array_real(i) &
      = derived_type%component_array_real(i) &
      + derived_type%component_array_real(derived_type%component_array_int(i))
  end do
end do
call system_clock(finish)
derived_type_time = derived_type_time + finish-start

call system_clock(start)
size = derived_type%component_length
dummy_array_int => derived_type%component_array_int
dummy_array_real => derived_type%component_array_real
do j = 1, 5
  dot = 0.0_rKind
  do i = 1, size
    dot = dot + dummy_array_real(i)*vector(i)
  end do
  max_element = -huge(max_element)
  do i = 1, size
    if ( abs(dummy_array_real(i)) > abs(max_element) ) &
      max_element = dummy_array_real(i)
  end do
  sum = 0.0_rKind
  do i = 1, size
    sum = sum + dummy_array_real(i)
  end do
  do i = 2, size
    dummy_array_real(i) &
      = dummy_array_real(i) &
      + dummy_array_real(i-1)
  end do
  do i = 1, size
    dummy_array_real(i) &
      = dummy_array_real(i) &
      + dummy_array_real(dummy_array_int(i))
  end do
end do
call system_clock(finish)
deref_derived_type_time = deref_derived_type_time + finish-start

end subroutine work_with_array_types

end module test_various_array_types

program test_array_storage_performance

use kind_definitions, only: iKind, rKind
use type_definition, only: derived
use module_data, only: module_data_array_int, module_data_array_real, &
  module_length

use test_various_array_types, only: read_array_types, &
  work_with_array_types, &
  reset_counters, &
  assumed_size_time, &
  assumed_shape_time, &
  module_data_time, &
  derived_type_time, &
  deref_derived_type_time

implicit none

integer, parameter :: logical_unit = 1

```

```

integer, parameter :: number_of_runs = 10
integer             :: array_size    = 20000

integer(iKind), dimension(:), allocatable, target :: assumed_size_array_int
real(rKind),    dimension(:), allocatable, target :: assumed_size_array_real

! integer(iKind), dimension(:), allocatable :: assumed_shape_array_int
! real(rKind),    dimension(:), allocatable :: assumed_shape_array_real

type(derived) :: derived_type

integer :: i, ii

continue

do ii=1, 3

    write(*,*) " timings for array length ",array_size

    allocate( assumed_size_array_int(array_size), &
              assumed_size_array_real(array_size) )
!   allocate( assumed_shape_array_int(array_size), &
!             assumed_shape_array_real(array_size) )

    module_length = array_size
    allocate( module_data_array_int(array_size), &
              module_data_array_real(array_size) )

!   allocate( derived_type%component_array_int(array_size), &
!             derived_type%component_array_real(array_size) )

    derived_type%component_length = array_size
    derived_type%component_array_int => assumed_size_array_int
    derived_type%component_array_real => assumed_size_array_real

    open (logical_unit, file='data', form='unformatted' )

    write (logical_unit) ((array_size-i+1), i=1,array_size)
    write (logical_unit) (1.0_rKind, i=1,array_size)

    do i = 1, number_of_runs
        call read_array_types( logical_unit, array_size,
                               &
                               assumed_size_array_int, assumed_size_array_real, &
                               derived_type%component_array_int, derived_type%component_array_real, &
                               assumed_shape_array_int, assumed_shape_array_real, &
                               derived_type%component_array_int, derived_type%component_array_real, &
                               derived_type )
    end do

    close (logical_unit)

    write(*,*) 'IO tests:'
    write(*,*) '    Assumed-size:', assumed_size_time
    write(*,*) '    Assumed-size:', assumed_size_time / real( assumed_size_time)
    write(*,*) '    Assumed-shape:', assumed_shape_time / real( assumed_size_time)
    write(*,*) '    Use Module:', module_data_time / real( assumed_size_time)
    write(*,*) '    Derived type:', derived_type_time / real( assumed_size_time)
    write(*,*) '    DeRef-Derived type:', deref_derived_type_time / real( assumed_size_time)

    call reset_counters()

    do i = 1, number_of_runs
        call work_with_array_types(array_size,
                                   &
                                   assumed_size_array_int, assumed_size_array_real, &
                                   derived_type%component_array_int, derived_type%component_array_real, &
                                   assumed_shape_array_int, assumed_shape_array_real, &
                                   derived_type%component_array_int, derived_type%component_array_real, &
                                   derived_type )
    end do

    write(*,*) 'Work tests:'
    write(*,*) '    Assumed-size:', assumed_size_time
    write(*,*) '    Assumed-size:', assumed_size_time / real( assumed_size_time)
    write(*,*) '    Assumed-shape:', assumed_shape_time / real( assumed_size_time)
    write(*,*) '    Use Module:', module_data_time / real( assumed_size_time)
    write(*,*) '    Derived type:', derived_type_time / real( assumed_size_time)
    write(*,*) '    DeRef-Derived type:', deref_derived_type_time / real( assumed_size_time)

    deallocate( assumed_size_array_int, assumed_size_array_real )
!   deallocate( assumed_shape_array_int, assumed_shape_array_real )

    deallocate( module_data_array_int, module_data_array_real)

!   deallocate( derived_type%component_array_int,
!             &
!             derived_type%component_array_real )

```

```
        array_size = array_size*10
    end do
end program test_array_storage_performance
```

